Subscribe to Home Delivery | bart.selman... | Help

The New York Times

Business Day
Technology

Search All NYTimes.com [_____] Go
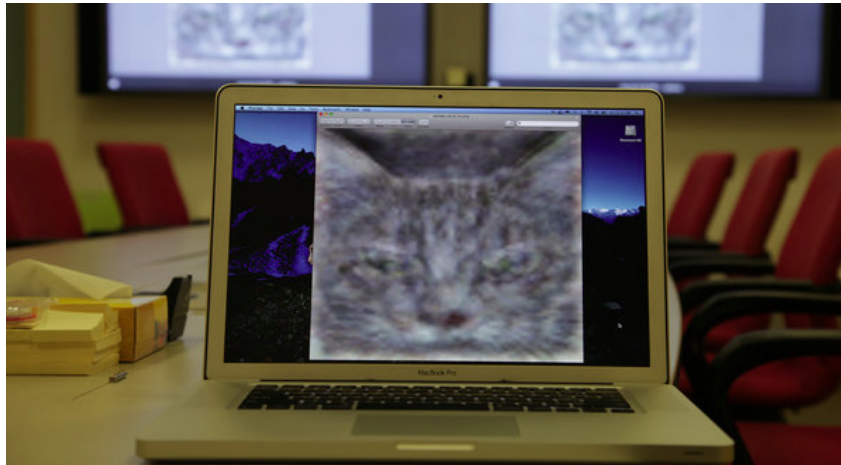
Capital One 360

How Many Computers to Identify a Cat? 16,000

Jim Wilson/The New York Times
An image of a cat that a neural network taught itself to recognize.

By JOHN MARKOFF
Published: June 25, 2012

MOUNTAIN VIEW, Calif. — Inside Google's secretive X laboratory, known for inventing self-driving cars and augmented reality glasses, a small group of researchers began working several years ago on a simulation of the human brain.

FACEBOOK
TWITTER
GOOGLE+
E-MAIL
SHARE
PRINT
REPRINTS

STOKER
COMING SOON

Multimedia
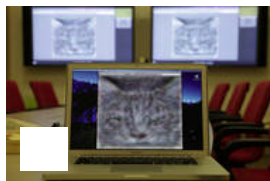
Business Day Live | Google's Brain

There Google scientists created one of the largest neural networks for machine learning by connecting 16,000 computer processors, which they turned loose on the Internet to learn on its own.

Presented with 10 million digital images found in YouTube videos, what did Google's brain do? What millions of humans do with YouTube: looked for cats.

Enlarge This Image

Jim Wilson/The New York Times
Andrew Y. Ng, a Stanford computer scientist, is cautiously optimistic about neural networks.

The neural network taught itself to recognize cats, which is actually no frivolous activity. This week the researchers will present the results of their work at a conference in Edinburgh, Scotland. The Google scientists and programmers will note that while it is hardly news that the Internet is full of cat videos, the simulation nevertheless surprised them. It performed far better than any previous effort by roughly doubling its accuracy in recognizing objects in a challenging list of 20,000 distinct items.

The research is representative of a new generation of

Log in to see what your friends are sharing on nytimes.com.
Privacy Policy | What's This?

Log In With Facebook

What's Popular Now f

A Statement Rocks Rome, Then Sends Shockwaves Around the World

Pope Benedict XVI Says He Will Retire

Today's Headlines Daily E-Mail

Subscribe to Technology RSS Feeds

Technology News	Bits Blog
Internet	Personal Tech
Business Computing	Pogue's Posts
Start-Ups	
Companies	

MOST E-MAILED   MOST VIEWED

1. WELL
Straining to Hear and Fend Off Dementia

2. OP-ED CONTRIBUTOR
Farewell to an Uninspiring Pope

3. OPINION
Relax! You'll Be More Productive

4. A Call for Drastic Changes in Educating New Lawyers

computer science that is exploiting the falling cost of computing and the availability of huge clusters of computers in giant data centers. It is leading to significant advances in areas as diverse as machine vision and perception, speech recognition and language translation.

Although some of the computer science ideas that the researchers are using are not new, the sheer scale of the software simulations is leading to learning systems that were not previously possible. And Google researchers are not alone in exploiting the techniques, which are referred to as "deep learning" models. Last year Microsoft scientists presented research showing that the techniques could be applied equally well to build computer systems to understand human speech.

"This is the hottest thing in the speech recognition field these days," said Yann LeCun, a computer scientist who specializes in machine learning at the Courant Institute of Mathematical Sciences at New York University.

And then, of course, there are the cats.

To find them, the Google research team, led by the Stanford University computer scientist Andrew Y. Ng and the Google fellow Jeff Dean, used an array of 16,000 processors to create a neural network with more than one billion connections. They then fed it random thumbnails of images, one each extracted from 10 million YouTube videos.

The videos were selected randomly and that in itself is an interesting comment on what interests humans in the Internet age. However, the research is also striking. That is because the software-based neural network created by the researchers appeared to closely mirror theories developed by biologists that suggest individual neurons are trained inside the brain to detect significant objects.

Currently much commercial machine vision technology is done by having humans "supervise" the learning process by labeling specific features. In the Google research, the machine was given no help in identifying features.

"The idea is that instead of having teams of researchers trying to find out how to find edges, you instead throw a ton of data at the algorithm and you let the data speak and have the software automatically learn from the data," Dr. Ng said.

"We never told it during the training, 'This is a cat,' " said Dr. Dean, who originally helped Google design the software that lets it easily break programs into many tasks that can be computed simultaneously. "It basically invented the concept of a cat. We probably have other ones that are side views of cats."

The Google brain assembled a dreamlike digital image of a cat by employing a hierarchy of memory locations to successively cull out general features after being exposed to millions of images. The scientists said, however, that it appeared they had developed a cybernetic cousin to what takes place in the brain's visual cortex.

Neuroscientists have discussed the possibility of what they call the "grandmother neuron," specialized cells in the brain that fire when they are exposed repeatedly or "trained" to recognize a particular face of an individual.

"You learn to identify a friend through repetition," said Gary Bradski, a neuroscientist at Industrial Perception, in Palo Alto, Calif.

While the scientists were struck by the parallel emergence of the cat images, as well as human faces and body parts in specific memory regions of their computer model, Dr. Ng said he was cautious about drawing parallels between his software system and biological life.

"A loose and frankly awful analogy is that our numerical parameters correspond to synapses," said Dr. Ng. He noted that one difference was that despite the immense computing capacity that the scientists used, it was still dwarfed by the number of connections found in the brain.

"It is worth noting that our network is still tiny compared to the human visual cortex, which is a million times larger in terms of the number of neurons and synapses," the researchers wrote.

Despite being dwarfed by the immense scale of biological brains, the Google research provides new evidence that existing machine learning algorithms improve greatly as the machines are given access to large pools of data.

"The Stanford/Google paper pushes the envelope on the size and scale of neural networks by an order of magnitude over previous efforts," said David A. Bader, executive director of high-performance computing at the Georgia Tech College of Computing. He said that rapid increases in computer technology would close the gap within a relatively short period of time: "The scale of modeling the full human visual cortex may be within reach before the end of the decade."

Google scientists said that the research project had now moved out of the Google X laboratory and was being pursued in the division that houses the company's search business and related services. Potential applications include improvements to image search, speech recognition and machine language translation.

Despite their success, the Google researchers remained cautious about whether they had hit upon the holy grail of machines that can teach themselves.

"It'd be fantastic if it turns out that all we need to do is take current algorithms and run them bigger, but my gut feeling is that we still don't quite have the right algorithm yet," said Dr. Ng.

This article has been revised to reflect the following correction:

Correction: June 25, 2012

An earlier version of this article misstated the size of the human visual cortex.

A version of this article appeared in print on June 26, 2012, on page B1 of the New York edition with the headline: How Many Computers to Identify a Cat? 16,000.

FACEBOOK   TWITTER   GOOGLE+   E-MAIL   SHARE

INSIDE NYTIMES.COM   ◀ ▶

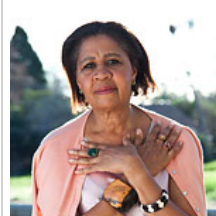| HEALTH »	BUSINESS »	OPINION »	DINING & WINE »	OPINION »	BOOKS »
A Few Workouts a Week May Be Better Than More	In Europe, a Risk-Filled Choice for Britain	Op-Ed: Germany's Political Masochism	Restaurant Review: Aska	**Can Danger Be Bred Into a Bloodline?** Room for Debate asks: Does selecting for traits that make a pedigree dog a champion create traits that can hurt it?	'See Now Then,' by Jamaica Kincaid